

CLAIMS

What is claimed is:

1. A method of status generation for a node of a high-availability cluster, the
5 method comprising:
sending a heartbeat signal from the node through a network to the cluster;
determining a current status of the node; and
sending the current status out through a specialized interface to a next
node.
- 10 2. The method of claim 1, wherein the specialized interface is dedicated to
inter-node status communication, and wherein the network is used for
other communications in addition to the heartbeat signaling.
- 15 3. The method of claim 1, further comprising:
determining a current degraded level of the node; and
sending the degraded level out through the specialized hardware to the
next node.
- 20 4. The method of claim 3, wherein the specialized interface is dedicated to
inter-node status communication, and wherein the network is used for
other communications in addition to the heartbeat signaling.
- 25 5. The method of claim 4, wherein the specialized interface couples nodes of
the cluster in a ring topology.
- 30 6. A method of cluster-wide management performed per node, the method
comprising:
checking an up/down status input received from a previous node;
checking a degraded status input received from the previous node; and
checking a heartbeat input received from the previous node.

7. The method of claim 6, wherein the degraded status input comprises multiple degradation levels, and wherein one such level comprises a "bad" state indicating that the previous node appears down.

5 8. The method of claim 6, further comprising:
determining whether a configuration file at the previous node has been
changed; and
if the configuration file has been changed, then retrieving the configuration
file from the previous node and storing the retrieved configuration
10 file at the present node.

9. The method of claim 6, further comprising:
performing a logical analysis of the inputs to determine whether a failure
of the previous node is indicated.

15 10. The method of claim 9, wherein the logical analysis comprises
determining a failure of the previous node if a majority of the status inputs
indicates that the previous node appears down.

20 11. The method of claim 9, wherein the logical analysis differentiates between
the failure of the previous node and a failure of an inter-node
communication channel.

25 12. The method of claim 11, wherein the logical analysis further differentiates
between a problem with a first inter-node communication channel and a
problem with a second inter-node communication channel.

30 13. The method of claim 12, wherein the first inter-node communication
channel comprises a point-to-point link dedicated for node status
information, and wherein the second inter-node communication channel
comprises a network for carrying heartbeat signals and other
communications.

14. The method of claim 7, further comprising reporting that a network carrying the heartbeat is down if the heartbeat is bad and the two status inputs are not both bad.
- 5 15. The method of claim 7, further comprising reporting a problem with an inter-node communication channel carrying the status inputs if the heartbeat is okay and one, but not both, of the two status inputs is bad.
- 10 16. The method of claim 7, further comprising comparing the degraded status with a node removal threshold for potential removal of the previous node from the cluster if the degraded status shows degradation above the threshold.
- 15 17. A system for of a high-availability cluster, the system comprising:
a general inter-node communication network that is configured to carry signals including heartbeat signals from the nodes; and
a separate inter-node communication channel for communicating node status signals.
- 20 18. The system of claim 17, wherein the node status signals includes an up/down status signal and a degraded status signal.
- 25 19. The system of claim 18, wherein the system is configured with a logical analysis procedure that differentiates between a failure of a node and a problem with inter-node communication.
20. The system of claim 19, wherein the logical analysis further differentiates between a problem with the general inter-node communication network and a problem with the separate inter-node communication channel.